

The Hyperscaler Funding Gap


Hyperscalers starting to see ROIC on capex, press further into future cash flows

Jordan Rohan · Clearmeadow Partners, LLC · July 31, 2026

INVESTMENT HIGHLIGHTS

- Estimates for 2027 hyperscaler capex rose \$114 billion in a single earnings season.** Built name by name from published sell-side models, the seven-name total for 2027 moves from \$1,207 billion on 6 July to \$1,321 billion today. The gross figure is larger still: \$204 billion of raises at Amazon, Meta, Alphabet and Nebius against \$90 billion of reductions, and \$79 billion of those reductions are definitional. Microsoft's \$42 billion decline is the datacenter lease reclassification, which management states leaves underlying investment unchanged, and Oracle's \$37 billion is calendarization. CoreWeave alone is genuinely spending less.
- Self-funding ends this year, on both denominators.** Rebuilt from four current published models, the four largest spenders lay out \$738 billion of capex in 2026 against \$741 billion of operating cash flow, precisely parity, then \$1,175 billion against \$938 billion in 2027, a ratio of 125%. Free cash flow runs from \$200 billion in 2025 to roughly zero in 2026 and negative \$236 billion in 2027. Measured against earnings the crossover appears in the same years: capex runs at 100% of adjusted EBITDA this year and 118% in 2027, with Amazon at 99% and 108%, Meta at 96% and 130% and Alphabet at 85% and 125%. Microsoft at 86% is the only major still spending materially less than it earns before depreciation.
- Financed spending was tested this week, in public.** Capex above operating cash flow comes off the balance sheet by definition, and third-party work puts outstanding AI-related debt near \$7 trillion by 2029. Situational Awareness LP, a fund built on this thesis, was liquidated by its prime brokers on 30 July: up 439% through June, levered four times, long CoreWeave, Nebius, Bloom Energy, Lumentum and Coherent, and carried from \$45 billion of assets to \$10 billion with Citadel taking the entire public book. Demand held throughout. The financing structure is what gave way.
- The spending is self-correcting, which is the most important structural fact here.** Capex is reset annually against observed returns, so either the return materializes and the programs continue or it fails to and they moderate. Both outcomes are acceptable. Four management teams reached this view independently, each with its own board, compensation structure, silicon strategy and shareholders' money at risk. The correction works least cleanly at CoreWeave and Nebius, which spend against signed leases and debt covenants.
- The core franchises are the mechanism, and three of the four have one.** AWS, Azure and Google Cloud already earn their cost of capital on conventional workloads and are growing 37%, 43% and 82%, which gives incremental AI capacity a paying use from delivery. Amazon frames breakeven inside three years on equipment with five to six year useful lives, inside shells with thirty year lives that support successive silicon generations. Oracle frames project-level ROIC in the high 20s. Equity markets are sorting on exactly that evidence, and Meta matters least to this analysis until it demonstrates it can sell to enterprises.

Exhibit 1. Alphabet, Microsoft, Meta and Amazon consume all of their operating cash flow in 2026 and 125% of it in 2027




Note to Exhibit 1: Alphabet, Microsoft, Meta and Amazon on a calendar-year basis. History is rebuilt from company filings; forecast years are the sum of current per-name published models. Sourcing and reconciliation in Notes on the Estimates.

THE BUILDOUT IS BROADENING, WHICH RAISES THE STAKES

The spend extends well past the four largest balance sheets. The seven-name total reaches roughly \$1.32 trillion in 2027, and the fastest growth is at Oracle, CoreWeave and Nebius, whose 2024 to 2027 capex growth runs well above the majors; Nebius alone compounds from roughly \$0.1 billion to \$31.1 billion. That matters because the neoclouds fund capacity with debt and lease it to a concentrated set of customers. Where the incremental spending is financed and the revenue behind it is customer-concentrated, the fragility of the whole rises faster than the aggregate ratio suggests.

The clearing of the Situational Awareness book removes a known seller from the names it held. Roughly \$20 billion of concentrated long exposure to AI infrastructure, power and data center equities changed hands over a matter of days, against a set of shares whose average daily volume absorbs that size slowly. Several of the affected names rose 15% to 25% on the session following the liquidation. Supply of that character arrives once and exhausts itself, and the technical low it establishes tends to hold while the fundamental questions in this note remain open. We would expect the names in that portfolio to trade well over the coming weeks for reasons unrelated to their return on invested capital, and we would read that strength as the absence of a seller. Buyers of a forced block price it to move, so the bounce measures the concession rather than a change in the earnings picture.

Exhibit 2. The buildout is broadening: hyperscaler and neocloud capex, 2024-2027



Note to Exhibit 2: the 2025 capex column is sourced to 10-K and 10-Q filings and totals \$464 billion, against the \$435 billion we published in July. Exhibit 2 adds Oracle, CoreWeave and Nebius to the four names in Exhibit 1 and therefore carries a larger total; both are correct on their own definitions. Microsoft is shown on the reported basis. Company guidance implies roughly \$15 billion of calendar 2026 spending sits outside reported capex following the lease reclassification, so the group total is understated by about that amount on a like-for-like basis with the other six names.

WHAT CHANGED SINCE JULY 6

Name	2026E	% EBITDA	2027E prior	2027E new	% EBITDA	Change	Basis
Alphabet	\$202.6B	85%	\$340B	\$364.2B	125%	+\$24B	JPM and Barclays blended
Amazon	\$220.5B	99%	\$232B	\$330.8B	108%	+\$99B	Barclays / JPM / MS blend
Microsoft	\$173.7B	78%	\$280B	\$238.4B	86%	-\$42B	MS and JPM blended
Meta	\$141.3B	96%	\$175B	\$241.2B	130%	+\$66B	Barclays / JPM / MS blend
Oracle	\$77.2B	224%	\$132B	\$94.6B	212%	-\$37B	JPM 10 Jun; calendarized
CoreWeave	\$35.6B	486%	\$32B	\$20.9B	184%	-\$11B	JPM; peaks in 2026
Nebius	\$21.7B	1,600%	\$16B	\$31.1B	502%	+\$15B	MS 13 May; pre-2Q26
Seven-name group	\$873B	100%	\$1,207B*	\$1,321B	118%	+\$114B	All seven refreshed

THE REAL QUESTION IS THE RETURN

A capital program of this size is rational if, and only if, the capacity earns a return above its cost. That is the test the market has begun to apply, and it is why the shares of the spenders no longer re-rate on capex announcements the way they once did. On the evidence available today, the number of operators who can clearly clear that bar is small.

- ◆ **The frontier labs are the demand side of this analysis.** A large share of AI revenue at both the hyperscalers and the neoclouds traces to a small number of private labs, so the contracts standing behind the financed capacity are worth what those labs' economics are worth. Anthropic has the clearest path, with enterprise demand compounding and cash burn falling steadily. OpenAI carries a heavier consumer load against the same compute costs, which makes its margin path the more

consequential unknown. Neither can be owned, and both set the terms on which this capital gets repaid.

- ◆ **Google is the best-positioned public name, and it is paying up for capacity it cannot build fast enough.** It designs its own accelerators, funds the buildout from a business that already earns its cost of capital, and has the model traction to give new capacity a paying use. Cloud growth accelerated again and the backlog now exceeds half a trillion dollars. Alongside that, Alphabet is renting outside GPU capacity at above-market rates, as a bridge for cloud customers and to free its own machines for training. The best-capitalized operator in the group, with its own silicon, is paying a premium to rent someone else's. Nothing states the supply position more plainly, and it means reported capex understates what is actually being spent on capacity.
- ◆ **Alphabet's estimates rose more than any name here, and its return on capital moves the other way.** The raise came with matching increases to cloud revenue, so it follows demand rather than cost. Return on invested capital still declines steadily through 2028, because the asset base compounds faster than the earnings it produces. That pattern holds across the group and it is worth separating from the project-level economics operators describe, which sit above corporate overhead and financing.
- ◆ **Microsoft has strengthened on this test.** Cloud growth accelerated and management attributes it to bringing capacity online faster and running the existing fleet harder, describing demand as running ahead of supply. Contracted future revenue rose sharply, the enterprise AI product moved from pilots to scaled adoption, and margins held through the build. Free cash flow stays positive across the fiscal years while converging toward zero, and on a calendar basis it turns negative in 2027, when the heaviest spending falls. The strongest balance sheet in the group reaches the same crossover as everyone else, one year later and from a better starting position.
- ◆ **Meta carries hyperscaler capital intensity without a hyperscaler franchise.** Capex rises sharply next year while revenue estimates hold, and the capacity being built earns close to nothing today against peers earning real revenue on every unit of it. Returns fall and debt climbs across the forecast. That is a consumer company spending like an infrastructure company, and it stays that way until Meta demonstrates it can sell capacity to enterprises. Analysts have started to size what that would be worth, and we take the name up separately.
- ◆ **Amazon belongs in the foreground, and the quarter moved it there.** Guidance went up again on higher component costs, and the spending runs ahead of earnings next year after a long stretch of the reverse. The offsetting evidence is the strongest in the group: cloud growth reaccelerated hard, the AI and custom silicon businesses each cleared a \$25 billion run rate, and segment margins held. Management also put a payback framework on the table, arguing the equipment pays for itself inside three years and the buildings house several successive generations of it. That is the most concrete answer any operator has given to the question this note asks.
- ◆ **Oracle has put the most specific number on the table.** Management frames project-level returns for its infrastructure business in the high 20s, against a reaffirmed margin framework and a backlog that keeps setting records. That is the clearest return claim any operator here has offered, and it describes steady-state project economics, which sit above corporate overhead and the cost of the financing. The funding runs alongside it: capex exceeds revenue next fiscal year, and Oracle plans a large combined debt and equity raise to cover the gap.
- ◆ **Oracle's prepayment structure is a genuine mitigant, and it concentrates the customer risk.** Much of the new business is prepaid or arrives with customer-supplied hardware, so the capacity is funded ahead of delivery rather than by debt. The same structure ties the return to a handful of counterparties, which is the neocloud exposure arrived at by a different route.
- ◆ **The neoclouds show what the gap looks like at the end of the chain.** CoreWeave's operating earnings cover a little over half its interest bill in 2027, against debt that keeps climbing and almost no equity beneath it. Nebius spends several times its operating cash flow, and a meaningful share of that cash flow is customer prepayment rather than earnings. Set both against the returns Oracle describes and the range within this group becomes the whole argument. CoreWeave's capex also peaks this year and falls in 2027, the only decline among the seven, which is the sequence our thesis anticipates.

EXHIBIT 3. WHAT WOULD DEMONSTRATE THE RETURN

Name	Demonstration	Evidence to date	Still outstanding
Alphabet	TPU cost advantage flows through to cloud margin	Cloud +82%, backlog \$514B, Cloud margin 36%, first TPU hardware revenue, Gemini API 22B tokens/min	ROIC falls 26.9% to 15.2% by 2028; capex 140% of EBITDA in 2028
Microsoft	Azure growth absorbs capacity at stable margin	Azure 43%, ~45% guided, supply-constrained; RPO \$678B up 84%; Copilot 30MM seats; op margin 45.1%; FCF positive through FY29	Calendar 2027 free cash flow turns negative
Amazon	Payback framework holds across silicon generations	AWS 37% from 28%, AI and custom silicon each above \$25B run rate, 37.9% segment margin	Multi-vintage utilization; FCF -\$55B by 2028
Oracle	High-20s project ROIC realized above overhead and financing	\$638B backlog, \$67B signed in Q4, prepayments fund \$20-25B	Corporate-level return; four customers above \$8B each
Meta	External monetization of merchant capacity	Analysts size a \$33.5B and \$112.6B revenue opportunity and \$4.19 of 2028 EPS	Monetization near \$0 per GW today; timing undisclosed
CoreWeave	Operating earnings cover the interest bill	Capex peaks in 2026 and falls 41% in 2027	\$2.5B EBIT against \$4.2B interest; ROCE 5.9%
Nebius	Contracted capacity converts to cash	Company guide raised to \$20-25B	Capex 298% of operating cash flow in 2027; coverage 54%

BOTTOM LINE

HOW IT GETS FUNDED. Capex is funded three ways: from operating cash flow, from debt, or from equity issuance. The first is fully committed at current spending levels, the second is expensive while private credit is being questioned, and the third is available case by case. Alphabet has combined equity and debt placement to cover the near-term gap, Microsoft's spending stays within reach of its cash generation, and Amazon can fund from a mix of cash flow and issuance.

BASE CASE. The spending resolves one of two ways, and **both are acceptable**. Either the return materializes and these programs continue, or it fails to and capex growth moderates. Capex is reset annually against observed returns, which makes this self-correcting rather than a commitment that runs to its conclusion regardless of the evidence. Four management teams reached the same view independently, each with its own board, its own compensation structure, its own silicon strategy and its own shareholders' money at risk. The weight of that judgment favors the return arriving. Our base case is that it does.

WHAT WOULD INTERRUPT IT. Two conditions could. Open-weight models compress the pricing power on inference, since capacity earns rate-card revenue only while the frontier stays proprietary enough to charge for. And credit markets may decline to fund the back half of these programs on current terms, which is the constraint that binds before ROIC does. Either would slow the spending, and slower spending is the correction working rather than the thesis failing.

HOW THE MARKET IS SORTING. Equity markets are already sorting the group on evidence, and the sorting is the investable observation. Amazon, Microsoft and Alphabet carry current monetization traction in AWS, Azure and Google Cloud, and their shares are treated accordingly. Meta is the least consequential of the four to this analysis, since a company without a cloud franchise has the furthest to travel before its capacity earns anything, and it stays there until it proves it can sell to enterprises. The setup in the shares is a separate question and it has improved: the margin drag from the buildout now sits in consensus numbers while the offtake and consumer agent revenue does not, which is an asymmetry worth its own note.

WHERE THE RISK SITS. Our discipline concentrates at the end of the chain. The hyperscalers can slow a program that stops earning. CoreWeave and Nebius spend against signed leases and debt covenants, so their capacity is committed whether or not the pricing holds. That is where a moderation in the buildout is felt first and hardest, and it is where we would ask for the ROIC evidence before the multiple is paid.

NOTES ON THE ESTIMATES

Method. Clearmeadow estimates for capex, operating cash flow, free cash flow and adjusted EBITDA for all companies named in this report are derived from published sell-side company models from Barclays, J.P. Morgan, Morgan Stanley and other firms. Where more than one post-quarter model is available we blend them equally; otherwise we carry the single current model. Exhibit 1 forecast years are the sum of four such per-name models and its history is sourced to 10-K and 10-Q filings, as is the 2025 capex column in Exhibit 2. Building those forecast years from current models rather than carrying our July series forward raised 2026 operating cash flow by roughly \$25 billion and moved the crossover from 2026 into 2027, the year our July note identified. The 2024 column in Exhibit 2 remains a mixed-basis comparison base rather than a set of filed figures, and growth rates measured from it should be read with that in mind. Where published models disagree materially on a single line, as two do on Alphabet's 2027 free cash flow, we disclose the range rather than resolve it. Nebius is the one name whose model predates its second-quarter report.

Definitions. Adjusted EBITDA is stated with stock compensation added back. That definition is not uniform across published models, and one model we read omits the add-back, which understates the denominator by roughly \$20 billion a year at Amazon; we test each model against the standard rather than assume it by publisher. Depreciation on this buildout compounds quickly, so capex measured against EBITDA improves mechanically as the spending matures, and we read the level alongside the capex growth rate.

Two names need a definitional note. Effective FY27 Microsoft extends the estimated useful life of datacenters from 15 to 25 years, which moves more future leases out of reported capex and into cost of revenue. Company guidance for calendar 2026 falls from roughly \$190 billion to \$175 billion on that basis with the underlying investment unchanged, so Microsoft's bar in Exhibit 2 measures a narrower quantity than the other six by roughly \$15 billion and its apparent deceleration is an accounting artifact. Microsoft calendar figures are converted from June fiscal years. Oracle reports on a May fiscal year, so no exact calendar figure exists; its capex is gross, and a meaningful share of the FY27 program is expected to be funded by customer prepayments.

Other sources. The Situational Awareness account draws on CNBC and Bloomberg reporting of 30 July 2026. Third-party estimates of cumulative AI infrastructure spending and outstanding AI-related debt are from SemiAnalysis, mid-2026. Illustrative of the Clearmeadow view; not investment advice. Clearmeadow, its principals, or clients may hold positions in the securities discussed.